


AD-A175 206


ON THE CHARACTERIZATION OF SPATIAL FILTERS FOR
POINT-TARGET DETECTION IN IR IMAGERY(U) DEFENCE
RESEARCH ESTABLISHMENT VALCARTIER (QUEBEC) L SEVIGNY
NOV 86 DREV-4421/86 F/G 17/5

1/1

UNCLASSIFIED

NL





PHOTOCOPY RESOLUTION TEST CHART



National Défense
Defence nationale

UNCLASSIFIED
UNLIMITED DISTRIBUTION

DREV REPORT 4421/86
FILE: 3633J-005
NOVEMBER 1986

CRDV RAPPORT 4421
DOSSIER: 3633J-005
NOVEMBRE 1986

ON THE CHARACTERIZATION OF SPATIAL FILTERS FOR
POINT-TARGET DETECTION IN IR IMAGERY

L. Sévigny

DTIC FILE COPY

AD-A175 206

DTIC
ELECTE
DEC 12 1986
B

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

RESEARCH AND DEVELOPMENT BRANCH
DEPARTMENT OF NATIONAL DEFENCE
CANADA

BUREAU - RECHERCHE ET DÉVELOPPEMENT
MINISTÈRE DE LA DÉFENSE NATIONALE
CANADA

Defence Research Establishment
Centre de recherches pour la Défense,
Valcartier, Québec

404945-5H

Canada

SANS CLASSIFICATION
DISTRIBUTION ILLIMITÉE

86 12 11 1

DREV R-4421/86
FILE: 3633J-005

UNCLASSIFIED

CRDV R-4421/86
DOSSIER: 3633J-005

ON THE CHARACTERIZATION OF SPATIAL FILTERS FOR
POINT-TARGET DETECTION IN IR IMAGERY

by

L. Sévigny

DEFENCE RESEARCH ESTABLISHMENT
CENTRE DE RECHERCHES POUR LA DÉFENSE

VALCARTIER

Tel: (418) 844-4271

Québec, Canada

November/novembre 1986

SANS CLASSIFICATION

UNCLASSIFIED


1

ABSTRACT

This report is a comparative study of the spatial filters proposed for detecting stationary or slow-moving point targets. The filters are characterized by their probability of detection against a) the probability of false alarm, b) the strength or contrast of the target for a given false-alarm rate or c) the noise for a given false-alarm rate and a given strength or contrast.

RÉSUMÉ

Ce rapport consiste en une étude comparative des différents filtres spatiaux qui sont proposés pour la détection de cibles ponctuelles stationnaires ou quasi stationnaires. Les filtres ont été caractérisés par leur probabilité de détection en fonction: a) de la probabilité d'une fausse alarme, b) de l'intensité ou du contraste de la cible pour une probabilité de fausse alarme donnée ou c) du bruit pour une probabilité de fausse alarme et une intensité ou un contraste donnés.



Accession For	
NTIS SPAGI	<input checked="checked" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution	
Availability Codes	
Dist	Avail and/or Special
A-1	

TABLE OF CONTENTS

ABSTRACT/RÉSUMÉ	1
1.0 INTRODUCTION	1
2.0 CHARACTERIZATION METHODOLOGY	3
2.1 Baseline Surveillance System	3
2.2 Probability of a False Alarm	6
2.3 Probability of a True Alarm	7
2.4 Filter Operating Characteristic	13
2.5 Implementation	20
3.0 SPATIAL FILTERS EVALUATED	20
4.0 PERFORMANCE EVALUATION RESULTS	27
5.0 CONCLUSION	41
6.0 REFERENCES	43

FIGURES 1 to 20

TABLES I to III

1.0 INTRODUCTION

Spatial filters are widely used for target-detection applications. Hence, during the past few years we have developed and tested a certain number of nonlinear spatial filters generically referred to as double-gated filters. Several other nonlinear and/or linear spatial filters have been developed for the same purpose by Defence Research Establishment Valcartier (DREV) or by private firms under contract from DREV (Refs. 1-2).

Given the large number of spatial filters that are discussed in the scientific literature, there is undoubtedly a need for some sort of characterization procedure that would allow one to assess the ability of a particular spatial filter to detect targets. We will limit ourselves here to the detection of point targets in cluttered IR backgrounds. This shall be construed as meaning that the characterization procedure to be developed is not applicable to the detection of extended targets and that the results obtained for point targets cannot be generalized for extended or spread targets. We indeed feel that point targets and extended targets should be treated separately. It is far from obvious that the same spatial filter can be equally good both for point targets and for extended targets. Most spatial filters designed to detect a point target are based on the principle that such a target is nothing but an intensity surge, and as such is characterized by high spatial-frequency contents as opposed to the background, which mainly consists of low spatial frequencies. The rationale is that by filtering out the low-frequency components of an image, one should enhance the target-to-clutter ratio and make the target easier to isolate from the background. An extended target, however, contains both high spatial frequencies (the edges of the target) and low spatial frequencies, i.e. the interior of the target. In consequence, by filtering out the low frequencies, one ends up with only the edges or the outline of the extended target, and the processing needed to iso-

late it is radically different from that needed to isolate a point target (a simple threshold) and is not amenable to the same characterization procedure.

In this report, we define a point target as a target whose image is fully contained in a single picture element, and an extended target as one whose image overlaps several picture elements or pixels. However, one should be aware that as the target moves, it is bound to be seen by more than one pixel, e.g. when it crosses the boundary between two pixels. This is because a point target has a finite extent and is not a point in the geometrical sense. Here we neglect this aspect of the problem. In other words, we assume that the target is always sitting on a single pixel and that if it moves, it does so in steps of an integral number of pixels. This is not really a restriction as far as the characterization procedure is concerned. Assume that the point target in question is square and that it fully fills a pixel in size. Then, as this target moves, part of it will necessarily be seen by more than one pixel. In fact, it will be seen by two pixels (intensity of the target split among two pixels) if it moves horizontally or vertically, and by four pixels if it moves at a certain angle. This means that the proportion of the target seen by a given pixel will vary from 1 to 0.25 as it moves. This also means, assuming we use a single-intensity threshold to detect the target, that it might not be detectable in some portions of its trajectory even though the target maneuvers against a uniform background.

In Chapter 2.0 we fully describe the proposed characterization methodology. For instance, we explain what we mean by a probability of a false alarm or a probability of a true alarm and how we proceed to calculate them. The next chapter briefly describes the spatial filters proposed to detect slow or stationary targets. Finally, Chapter 4.0 presents and discusses the characterization results obtained.

This work was performed at DREV between April and December 1985 under PCN 33J05, Signal Processing - SBIRS.

2.0 CHARACTERIZATION METHODOLOGY

2.1 Baseline Surveillance System

Figure 1 illustrates the method of operation of a baseline surveillance system consisting of (1) an IR sensor, (2) a spatial filter, and (3) a threshold device. The image gathered by the IR sensor is first filtered using, as in the example here, the following 3 x 3, linear high-pass spatial filter:

$$\begin{bmatrix} +0.25 & -0.50 & +0.25 \\ -0.50 & +1.00 & -0.50 \\ +0.25 & -0.50 & +0.25 \end{bmatrix} \quad [1]$$

The image is then thresholded according to a single intensity threshold. The resulting image consists of a number of scattered white spots that mark the location of those pixels in the original image whose intensity exceeds the threshold. We will refer here to the threshold exceedances as alarms and distinguish between true alarms, i.e. those threshold exceedances that correspond to real targets, and false alarms, i.e. those that arise from background clutter. We will use the word "detections" to designate the true alarms.

The above-mentioned filtering and thresholding operations constitute only the preliminary phase, what I.W. Kay (Ref. 3) calls the detection phase, of the decision-making process aimed at discriminating between targets and backgrounds within the sensor's field of search. Tracking algorithms, which distinguish between the resulting true and false alarms by means of their supposedly different characteristics observed over time, make the final decision whether or not to declare that a target is present. This second phase is referred to as the




FIGURE 1 - This figure illustrates the method of operation of a baseline surveillance system consisting of: (1) an IR sensor, (2) a high-pass spatial filter, and (3) a threshold device. The sequence on the left-hand side depicts the output of the system in the absence of targets, whereas the sequence on the right-hand side shows what the output might look like when targets are present. The original image was collected by the Netherlands IRSCAN system in the 8-12 μm band. This IR sensor consists of a vertical array of 120 cadmium mercury telluride detectors. They produce a complete image of 220 x 120 pixels over the horizontal scanning of a scene.

declaration phase in Ref. 3. The scope of this report is restricted to the detection phase and its primary objective is to formulate a methodology for assessing the ability of a particular linear or nonlinear spatial filter to detect point targets.

The detection phase is of the utmost importance because it should discard as much background as possible without throwing away a valid target. The critical nature of this phase is an accepted fact as noted in image-processing literature, where the term "segmentation" is used to designate it. On the one hand, the detection phase sets a ceiling on the probability of target detection, for a target mistakenly rejected at this step is lost for good. On the other hand, as pointed out in Ref. 3, tracking algorithms used in the subsequent declaration phase will work only if, initially, the expected number of false alarms is below a certain critical value. Moreover, the effectiveness of a tracking algorithm is extremely sensitive to errors unless the a priori false-alarm probability can be reduced by the signal-processing techniques peculiar to the detection phase (Ref. 3).

It is common for a surveillance system to be used under conditions where most of the time there are no targets at all within the sensor's field of search. Indeed, we expect a surveillance system to give the alert only in rare and exceptional occasions. This means that the vast majority of the alarms generated by the detection phase are false ones. On the other hand, as mentioned in the preceding paragraph, the results of the declaration phase depend heavily on the number of false alarms generated by the detection phase. Thus, it is important to set a limit on the number of false alarms. This is equivalent to prescribing a constant false alarm rate (CFAR). The number of alarms generated by the baseline surveillance system of Fig. 1 can be easily controlled by varying the level of the threshold. So, independent of the particular spatial filter used, it is always possible to select the threshold in such a way that the number of alarms will equal a preset number. Under these conditions, we

obviously cannot characterize a spatial filter by its ability to reduce the dispersion of the gray levels of the original image, as is sometimes done. Instead, given the CFAR condition, the problem of characterizing a particular spatial filter can be stated as follows. Among all the possible spatial filters used for detecting the presence of a point target with a given false-alarm probability, find the filter for which the probability of detection is a maximum. This is the well-known Neyman-Pearson criterion of detection theory (Ref. 4).

2.2 Probability of a False Alarm

An image filtered with the aid of a high-pass filter contains both positive and negative intensities. This implies that the threshold level can either be negative, to detect targets colder than the background; positive, to detect targets hotter than the background; or bipolar, to detect both types of targets. Let $n(t)$ be the number of alarms generated by a particular spatial filter, as a function of threshold level t , and let N be the size (total number of pixels) of the filtered image (not necessarily the same size as the original image, due to edge effects). Then, the probability of a false alarm, $P_f(t)$, is given by:

$$P_f(t-) = n(t-) / N = F_f(t-) \quad [2]$$

for a negative threshold, by:

$$P_f(t+) = 1 - n(t+) / N = 1 - F_f(t+) \quad [3]$$

for a positive threshold, and by:

$$P_f(t-, t+) = P_f(t-) + P_f(t+) \quad [4]$$

for a bipolar threshold.

In the above expressions $P_f(t-)$ is the cumulative distribution function (c.d.f.) of the filtered image. Note that, in the case of the bipolar threshold, if $|t-| = |t+| = t$ then $P_f(t-, t+) = P_f(t)$, i.e. the c.d.f. of the absolute value of the filtered image. Given a required $P_f(t)$ value, the CFAR threshold can be easily determined by calculating the percentile that divides the range of filtered gray-level values into the appropriate proportions.

In Fig. 1, the filtered gray-level values range from -93 to +93. However, for display purposes, the filtered image was normalized between -10 and +10. That is, full black corresponds to filtered intensities less than or equal to -10, whereas full white corresponds to filtered intensities greater than or equal to +10. The filtered image was thresholded at the 4.71 value in both cases (with and without targets). This particular value is the 99th percentile of the filtered image. In other words, the probability of a false alarm is 0.01. The original image consists of 120 rows of 220 pixels, whereas there are only 118 rows of 218 pixels in the filtered image. This is because the spatial filter is restricted within the confines of the image and under these conditions it is blind to the pixels sitting on the perimeter of the original image.

2.3 Probability of a True Alarm

We will now address the problem of determining the probability of a true alarm, i.e. the probability of detecting a target that suddenly appears somewhere within the sensor's field of search. To this end, we will make a few assumptions.

- a) Firstly, we will assume that the IR sensor is operated at its limit of resolution and that, in consequence, the target is not larger than the field of view (IFOV) of an individual pixel. In other words, as already mentioned, we assume that the target of interest is a point target.

- b) Secondly, we will assume that in order to establish a contact with a point target, it must appear on a single pixel. In practice, of course, a point target may at times be sensed by more than one pixel. In fact, if there is no gap between the individual IFOVs, the intensity of a point target may be shared by up to four pixels, which considerably reduces its detectability. This implies that we have to accept the fact that, with time, there might be breaks in the detection log, even though the target is maneuvering against a perfectly uniform background.
- c) Finally, we will follow the example of Ref. 3 and assume that the surveillance system is deliberately designed to announce that a true alarm has occurred if, and only if, the target coincides with the central pixel of the spatial filter. Otherwise, the target will be regarded as absent. The reason for this last assumption will become more obvious as we go along.

To determine the probability of a true alarm we will have to simulate, for lack of sufficient experimental data, the change that the image gathered by an IR sensor undergoes due to the presence of a point target within its field of search. Let us consider the IR cloud picture of Fig. 2. We prefer this picture here, for illustration purposes, to the one of Fig. 1 because the latter contains large, saturated areas. Assume that a point target suddenly crosses the sensor's field of search. Then as the target passes through the individual fields of view, the intensity of the corresponding pixels fluctuates. The amplitude of this fluctuation essentially depends on the intensity and the size of the target. However, it is not important here to know the exact mechanism behind this phenomenon. We will only retain that the intensity of a particular pixel is not the same with and without a




FIGURE 2 - The IR digital cloud picture, used to describe the proposed method of characterization of spatial filters, and its histogram. The horizontal axis of the graph ranges from 0 to 237, i.e. from the minimum gray-level value to the maximum gray-level value found in the picture, whereas the vertical axis ranges from 0 to the value of the mode of the histogram.

target in its IFOV and, hence, we will attempt to characterize a given spatial filter for a broad range of intensity variations, in a variety of scenarios.

The following four target-insertion scenarios cover many possibilities:

$$R(x, y) = K1 \quad [5]$$

$$R(x, y) = Ml(x, y) + K2 \times Ml(x, y) \quad [6]$$

$$R(x, y) = I(x, y) + K3 \quad [7]$$

$$R(x, y) = I(x, y) + K4 \times I(x, y) \quad [8]$$

where $R(x, y)$ is the intensity resulting from the presence of a point target within the IFOV of the pixel at spatial coordinate (x, y) ; $K1$, $K2$, $K3$ and $K4$ are constants; $I(x, y)$ is the intensity of the background, i.e. of the pixel in the absence of a target, and $Ml(x, y)$ is a local mean calculated over a neighborhood of the same size as the spatial filter but excluding the central pixel. The first scenario corresponds to a point target whose size is comparable to the IFOV, with the result that the intensity of the target is substituted for the intensity of the background. For performance evaluation purposes, constant $K1$ can be expressed in terms of the background mean: $k'Mg$; background standard deviation: $k''Sg$; or both, $Mg + k'''Sg$. It is worthy of note that in this scenario the target can be either colder or hotter than the immediate surroundings depending on the value of $K1$ and where the target is inserted. This means that to detect it one must use a bipolar threshold. Otherwise, the characterization results are meaningless. The other three scenarios can be associated with a point target smaller than the IFOV so that some of the background intensity combines with that of the target and must be accounted for. Equation 6 can be rearranged as follows:

$$K2 = \frac{R(x, y) - Ml(x, y)}{Ml(x, y)} \quad [9]$$

that is, K2 can be assimilated to local contrast. In the case of [8], the same rearrangement will show that K4 is a signal-to-noise ratio. Constant K3, like K1, can be expressed in terms of Mg, Sg or both. The point targets simulated with the aid of [6], [7] or [8] are either colder (negative K's) or hotter (positive K's) than the background, and must be detected accordingly so as not to bias the characterization results.

The graph in Fig. 2 is the histogram of the underlying cloud picture and an estimate of the probability density function (p.d.f.) of the background pixels. The horizontal axis of the graph ranges from 0 to 237, i.e. from the minimum to the maximum gray level present in the picture, whereas the vertical axis ranges from 0 to the value of the mode of the histogram. Both scales are linear. The very same histogram is repeatedly plotted on a semilogarithmic scale in Fig. 3 as a solid gray plot. The superimposed white curves in the same figure are the p.d.f. of the simulated point targets. The labels A, B, C, and D correspond to the scenarios [5], [6], [7] and [8] respectively. These curves were obtained by first inserting, according to the related scenario, a point target in each and every pixel of the picture of Fig. 2, and then calculating the histogram of the resulting collection of target-in-pixel intensities. However, as far as scenarios [5] and [7] are concerned, this procedure is not really necessary since the outcome is obvious. In the first case, the histogram is nothing but a spike (Fig. 3a) whose position on the horizontal axis is governed by constant K1. In the case of [7], the target-in-pixel histogram is a true replica of the background histogram but is shifted to the right or to the left, according to the sign of K3, by an amount equal to the absolute value of K3 (Fig. 3c). The other two histograms (Figs. 3b and d), although distinct, exhibit the same general trend: they both stretch the background histogram in one direction.




FIGURE 3 - The repeated gray curve is the histogram of the cloud background of Fig. 2 plotted on a semilogarithmic scale, whereas the superimposed white curves are the histograms of the inserted point targets corresponding to the four target-insertion scenarios described in Section 2.3.

Figure 4 shows the spreading of the background and target-in-pixel intensities after filtering with the spatial filter defined by [1]. In other words, Figs. 3 and 4 display the same quantities before and after filtering. The horizontal axis of all the plots of Fig. 4 ranges from -100 to +100. We note that the histogram of the filtered cloud picture is quite narrow, with a standard deviation of 2.27 compared to a standard deviation of 45.02 for the original, unfiltered picture. The histograms of the target-in-pixel intensities are much broader, except for the one in Fig. 4c. Since the filter used is linear, the target-in-pixel histogram associated with [7], as in Fig 3c, is simply a shifted version of the background histogram. The scenario that produces the broadest histogram is the one defined by [5] (Fig. 4a). It stretches out from -35 to 202.5 which explains the odd behavior of the curve at the positive end of the horizontal axis. On the other hand, the distinct nature of scenarios [6] and [8] stands out more clearly here (Figs. 4b and d) than in the case of Fig. 3.

As mentioned in Section 2.2, the histogram of a filtered image provides an estimate of the p.d.f. of a false alarm. In the same way, the histogram of the filtered target-in-pixel intensities provides an estimate of the p.d.f. of a true alarm for a given background scene. The probability per se of a true alarm or of a false one is obtained by cumulating their p.d.f. backwards. The resulting curves, Fig. 5, are never-increasing functions of the threshold level.

2.4 Filter Operating Characteristic

We can use the plots of Fig. 5 to determine the probability of detection (probability of a true alarm) and the probability of a false alarm associated with a particular threshold setting or, more realistically, to determine the probability of detection we can expect when we have to compromise with a certain false-alarm probability. Another way




FIGURE 4 - This figure is the counterpart of Fig. 3. It depicts the histograms of the cloud background and of the inserted point targets after filtering, whereas Fig. 3 displays the same quantities, but before filtering.




FIGURE 5 - By cumulating backwards the histogram of the cloud background and the histograms of the inserted point targets of Fig. 4, one obtains the probability of a false alarm (repeated gray curve), and the probability of detection (white curves) as a function of the threshold level. As in the case of Figs. 3 and 4, the letters A, B, C and D correspond to the four target-insertion scenarios described in Section 2.3.

of presenting these quantities is to plot the probability of detection versus the probability of a false alarm, for various target-contrast values (Fig. 6). Van Trees (Ref. 4) refers to such plots as the receiver operating characteristic. Here the expression filter operating characteristic is probably more appropriate. Another equivalent representation of these quantities consists in plotting the detection probability against target-contrast values, for a few selected false-alarm probabilities (Fig. 7). This last rearrangement of the data is the one we prefer for filter-characterization purposes. The false-alarm probability values selected are 0.001, 0.01, and 0.1. Given the size of the images considered here and the fact that we are dealing with single frames, it would not be realistic to select a probability of false alarm smaller than 0.001. On the other hand, since what we eventually want is a surveillance system having a high detection probability coupled with a low false-alarm rate, it is of no use to consider false-alarm probability values greater than 0.1.

Figure 7 shows that for a target-contrast value of 0.09 the detection probability is close to one, whatever the probability of a false alarm. In the presence of a degrading factor such as noise, the performance of a particular filter is bound to decline. To assess its performance under such conditions one can calculate, for a given target-contrast value, say 0.09, the corresponding detection probability for increasing amounts of simulated noise. The resulting curves are shown in Fig. 8. The additive, uncorrelated Gaussian noise was derived from a pseudo-random number generator. The noise had a zero-mean value and a standard deviation given by $N_f \times S_g$, where N_f is the noise factor and S_g the standard deviation of the whole image.

The plots of Figs. 7 and 8 can be reduced to a single number, a kind of figure of merit, for easier filter characterization and comparison. This number is referred to as the Minimum Detectable Contrast




FIGURE 6 - In Fig. 5, the probability of a false alarm and the probability of detection are each plotted against the threshold level for a particular value of the parameter characterizing a given target-insertion scenario. Another, more useful way of presenting these quantities is to plot directly the probability of detection versus the probability of a false alarm. This is what we have done here for the case of the target-insertion scenario described by [6] for the following values of K_2 : 0.01, 0.02 and 0.04.




FIGURE 7 - This figure is a variant of Fig. 6 where the probability of detection is plotted against the characteristic parameter of the target-insertion scenario for a few selected values of the probability of a false alarm, namely 0.001, 0.01 and 0.1.




FIGURE 8 - Behavior of the detection probability as more and more uncorrelated Gaussian noise is added to the system. The target-insertion scenario here is the same as the one used to generate Fig. 7, with the value of K_2 set at 0.09. The noise factor is the standard deviation of the noise divided by the standard deviation of the original image. The values of the detection probability at zero noise factor (Fig. 8) and at 0.09 target contrast (Fig. 7) are the same. Hence, the plots of Figs. 7 and 8 could be joined end to end.

in the case of Fig. 7, and as the Maximum Tolerable Noise in the case of Fig. 8. Both are simply the target-contrast value and the noise-factor value, corresponding to a detection probability of 50%.

2.5 Implementation

Figure 9 diagrammatically describes how the characterization procedure expounded in the previous sections has been implemented. This figure is largely self-explanatory. First, a moving window extracts an image sample having the same dimensions as the filter point-spread function (or the filter-defining equation in the case of a nonlinear filter). A point target is then added to the center pixel according to one of the scenarios defined in Section 2.3. Correlated or uncorrelated Gaussian noise is added next if the intent is to evaluate the performance of the filter in the presence of noise. Otherwise, this step is merely bypassed. The last operation is the filtering operation that simply consists (in the case of a linear filter) in multiplying the elements of the point-spread function by the corresponding elements of the modified moving window, and then in adding all these products together to end up with a single number (here 47.75) for each location of the moving window.

3.0 SPATIAL FILTERS EVALUATED

Some of the spatial filters proposed to detect slow or stationary targets are the:

a) 3 x 3-element, 4-neighbor Laplacian filter (F1)

This high-pass, linear spatial filter is the discrete counterpart of the analog Laplacian operator, i.e. of a second-order, two-dimensional derivative. Its point-spread function is given by:

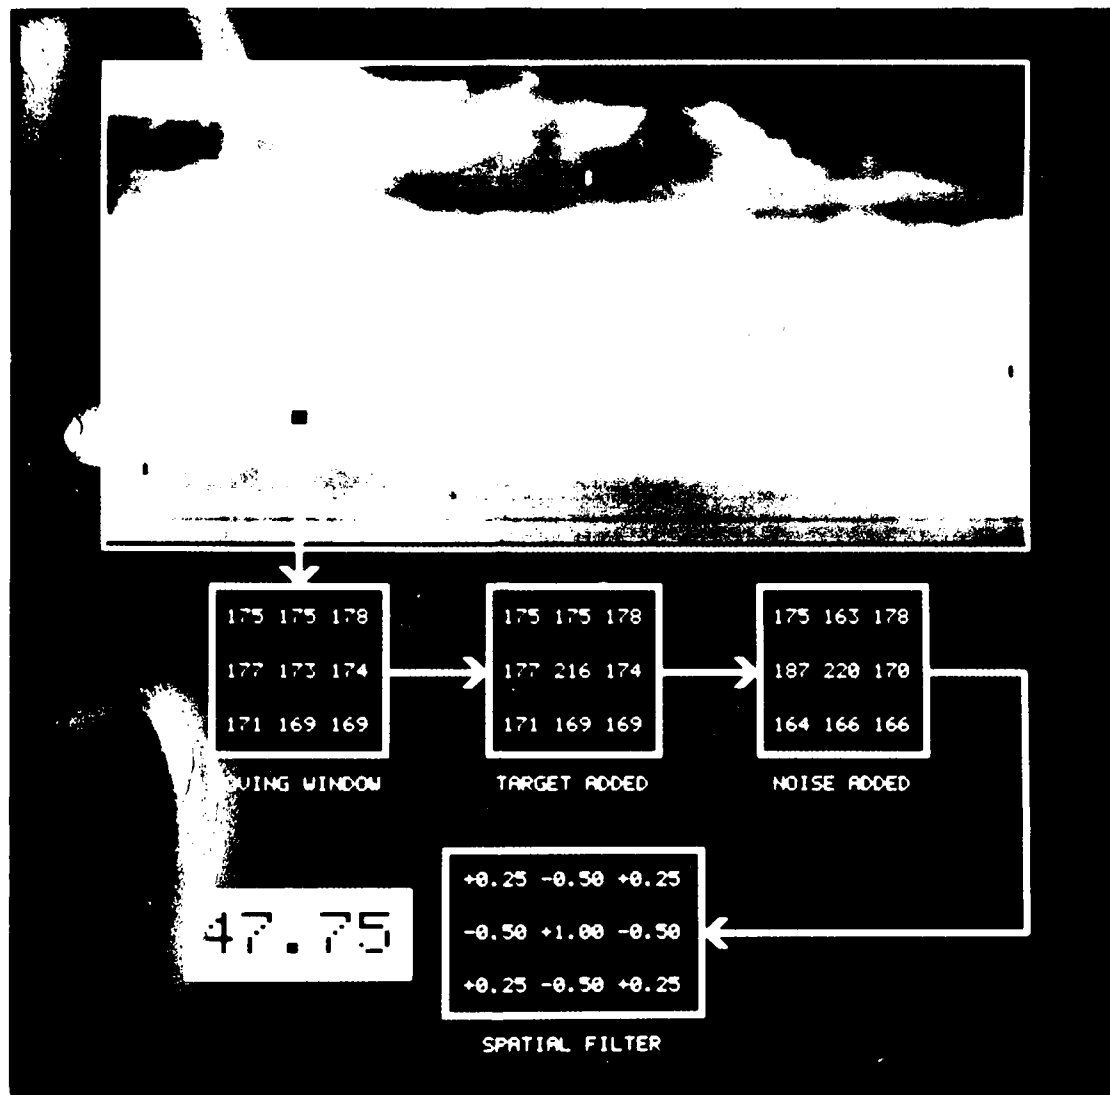


FIGURE 9 - This figure is largely self-explanatory. It describes the way the characterization procedure was implemented.

0.00	-0.25	0.00	
-0.25	+1.00	-0.25	[10]
0.00	-0.25	0.00	

b) 3 x 3-element, 8-neighbor Laplacian filter (F2)

This filter is an alternative Laplacian one based on a different neighborhood consisting of 8 horizontal, vertical, and diagonal neighbors weighted as follows:

+0.25	-0.50	+0.25	
-0.50	+1.00	-0.50	[11]
+0.25	-0.50	+0.25	

c) 5 x 5-element Wiener filter (F3)

0.0012	0.0195	0.0450	0.0196	0.0012	
0.0195	-0.0304	-0.3049	-0.0304	0.0195	
0.0450	-0.3049	1.0000	-0.3049	0.0450	[12]
0.0195	-0.0304	-0.3049	-0.0304	0.0195	
0.0012	0.0195	0.0450	0.0195	0.0012	

d) 9 x 9-element Wiener filter (F4)

0.0017	-0.0057	0.0091	-0.0314	0.0944	-0.0314	0.0091	-0.0057	0.0017	
-0.0057	0.0073	-0.0188	0.0520	-0.1599	0.0520	-0.0188	0.0073	-0.0057	
0.0092	-0.0188	0.0372	-0.1125	0.2954	-0.1125	0.0372	-0.0188	0.0092	
-0.0314	0.0520	-0.1125	0.2810	-0.5927	0.2810	-0.1125	0.0520	-0.0314	
0.0944	-0.1599	0.2954	-0.5927	1.0000	-0.5927	0.2954	-0.1599	0.0944	[13]
-0.0314	0.0520	-0.1125	0.2810	-0.5927	0.2810	-0.1125	0.0520	-0.0314	
0.0092	-0.0188	0.0372	-0.1125	0.2954	-0.1125	0.0372	-0.0188	0.0092	
-0.0057	0.0073	-0.0188	0.0520	-0.1599	0.0520	-0.0188	0.0073	-0.0057	
0.0017	-0.0057	0.0091	-0.0314	-0.0314	-0.0314	0.0091	-0.0057	0.0017	

Filters F3 and F4 are different least-mean-square optimal filters. The formulation of the optimal filtering problem requires a knowledge of the target and the background models. Here, we assume that we have to deal with a point target embedded respectively in a benign (F3) and in a harsh background (F4). The power spectral density of the background is assumed to be an exponential, the exponent of which determines the harshness of the background. These particular Wiener filters were proposed by E. Lithopoulos of Spar Aerospace Limited.

- e) 3 x 3-element submedian filter (F5)
- f) 5 x 5-element submedian filter (F6)

A one-dimensional median filter is defined by an odd-length window $L = 2*1 + 1$ elements long. To derive the output sample $g(i)$, the window is centered on the input sample $x(i)$ and all the samples within the window are ranked according to their amplitude. The output, $y(i)$, is selected to be the $(1 + 1)$ -ranked (i.e. the median) input sample.

The median operator preserves constant backgrounds, slopes and edges. It also suppresses isolated pulses shorter than or equal to 1 pixel. By performing a simple subtraction of the median input, $x(i)$, from the median output, $y(i)$, we derive a dual operator called the submedian. Unlike the median filter, the L -element long, submedian filter suppresses constant backgrounds, slopes and edges, and preserves pulses which are less than or equal to 1 element long.

For the present two-dimensional F5 and F6 filters, we use two different submedian operators. One is 3 elements long (F5) and the other is 5 elements long (F6). The operator is applied in four directions in the following sequence:

4 1 3

2 * 2

3 1 4

Specifically, the one-dimensional submedian operator is first applied in the north-south (1) direction. The resulting intermediate matrix is then processed by the submedian operator in the east-west (2) direction. The procedure is then repeated in the southwest-northeast (3) and the northwest-southeast (4) directions.

The performance of the two-dimensional submedian operator depends on the window size and the separability of the operator. Simulations indicate that the 3 x 3-element operator performs best for point targets. However, in the case of leakage of the target to neighboring elements, significant target attenuation can result. For this reason, a window 5 elements long is probably better for extracting extended targets.

g) 5 x 5-element double-gated filter (F7)

This nonlinear filter is fully documented in classified DREV reports where it is used to detect extended vehicle targets in IR imagery. The particular form of this filter, proposed to detect slow or stationary targets, can be implemented as follows. Each point in the input image is surrounded by a

frame one pixel thick, and 5 pixels wide. The mean, M_f , and the standard deviation, S_f , of the intensity of the pixels in the frame are then computed and these quantities are used to shift and scale the intensity (A) of the center element of the 5 x 5 window, as follows:

$$A' = (A - M_f) / S_f \quad [14]$$

Figure 10 illustrates the working principle of a double-gated filter.

Another filter proposed to detect slow or stationary targets is the Sobel edge detector. This filter is to be implemented in two stages. The first stage is the generation of two intermediate data matrices from the original image by using the two Sobel convolution operators, S_1 and S_2 , defined by:

$$\begin{array}{ccc} -0.5 & 0.0 & +5.0 \\ S_1 = -1.0 & 0.0 & +1.0 \\ -0.5 & 0.0 & +5.0 \end{array} \quad \begin{array}{ccc} -0.5 & -1.0 & -0.5 \\ S_2 = 0.0 & 0.0 & 0.0 \\ 0.5 & +1.0 & +0.5 \end{array} \quad [15]$$

We immediately see that such a filter cannot be evaluated by using the procedure expounded here, for the center element of both S_1 and S_2 is 0.

DOUBLE-GATED FILTER




FIGURE 10 - This figure illustrates the working principle of a double-gated filter. The black and the white disks superimposed on the target correspond respectively to the center of the white and gray frames surrounding the target. The white (black) disk and the white (gray) frame stick together and as the disk scans the target, the frame samples the background surrounding the target. For this reason, the disk or the inner window of a double-gated filter is referred to as the target gate, and the frame-like window or outer window as the background gate.

4.0 PERFORMANCE EVALUATION RESULTS

Figures 11 to 17 display the characteristic curves of the seven spatial filters defined in the previous chapter. The target-insertion scenario used to generate these curves is the one defined by [7]. The upper graphs in the aforementioned figures give the probability of detection as a function of the intensity of the target (constant K_3 of [7]), whereas the lower graphs depict the behavior of the probability of detection in the presence of noise. In the latter case, constant K_3 was set to 11 and the noise model used was:

$$N(0, N_f \times I(x, y)) \quad [16]$$

i.e. additive Gaussian noise with a mean of 0 and a standard deviation proportional to the intensity of the pixel at spatial coordinate (x, y) . N_f in [16] is the noise factor of Figs. 11 to 17.

From the plots of Figs. 11 to 17, one can extract the two figures of merit defined in Section 2.4, namely the minimum detectable target intensity (MDTI) and the maximum tolerable noise factor (MTNF). The values of these figures of merit are gathered in Tables I and II.

TABLE I

Minimum detectable target intensity

Filter	Probability of false alarm		
	0.001	0.01	0.1
4-neighbor Laplacian	>11.0	8.27	2.81
8-neighbor Laplacian	6.37	4.23	2.25
5 x 5 Wiener	>11.0	8.43	2.82
9 x 9 Wiener	>11.0	7.50	3.13
3 x 3 submedian	>11.0	7.47	3.52
5 x 5 submedian	>11.0	9.15	4.15
double-gated	>11.0	9.66	3.70




FIGURE 11 - The upper graph in this figure represents the probability of detecting a point target embedded in the cloud background of Fig. 2 by using a 3×3 , 4-neighbor Laplacian filter as a function of the intensity of the target above the intensity of the background, at the target location. The lower graph depicts the adverse effect of additive Gaussian noise on the probability of detecting a point target whose intensity is 11 units above the background. The noise in question has a mean of 0 and a standard deviation proportional (noise factor) to the intensity of each pixel within the filter window. Each graph displays three curves corresponding to the indicated false-alarm probabilities.




FIGURE 12 - This figure is the counterpart of Fig. 11 for the case of a 3×3 , 8-neighbor Laplacian filter. This filter is the one that exhibits the best overall performance for the cloud background considered here.




FIGURE 13 - This figure is the counterpart of Fig. 11 for the case of a 5 x 5 Wiener filter. The point-spread function of this filter is similar to the one of the 3 x 3, 4-neighbor Laplacian filter. This explains why the above plots are nearly a replica of those of Fig. 11.




FIGURE 14 - This figure is the counterpart of Fig. 11 for the case of a 9 x 9 Wiener filter.




FIGURE 15 - This figure is the counterpart of Fig. 11 for the case of a 3 x 3 submedian filter. This filter is a nonlinear filter, whereas the filters of Figs. 11 to 14 are linear.




FIGURE 16 - This figure is the counterpart of Fig. 11 for the case of a 5 x 5 submedian filter. As we can see by comparing the plots of Figs. 15 and 16, a larger submedian filter does not necessarily mean a better probability of detection.




FIGURE 17 - This figure is the counterpart of Fig. 11 for the case of a 5 x 5 double-gated filter. This filter, like the submedian filter, is a nonlinear filter. Although its overall detection performance is not impressive, it performs better than most of the other filters at low target intensities and it is more robust in the presence of noise.

TABLE IIMaximum tolerable noise factor

Filter	Probability of false alarm		
	0.001	0.01	0.1
4-neighbor Laplacian	0.0	0.013	0.047
8-neighbor Laplacian	0.013	0.017	0.040
5 x 5 Wiener	0.0	0.012	0.045
9 x 9 Wiener	0.0	0.014	0.031
3 x 3 submedian	0.0	0.015	0.044
5 x 5 submedian	0.0	0.011	0.043
double-gated	0.0	0.006	0.049

The next table (Table III) contains a few statistical parameters pertaining to the filtered image, i.e. the mean value, M_f , the standard deviation, S_f , and the AC power gain, G . The latter is defined as follows:

$$G = 10 \log_{10} (S_g^2 / S_f^2) \quad [17]$$

where S_g is the standard deviation of the original, unfiltered image.

TABLE IIIFiltered-image statistics

Filter	Mean (Mf)	Standard deviation (Sf)	AC power gain (G)
4-neighbor Laplacian	0.00	5.74	17.89
8-neighbor Laplacian	0.00	2.30	25.83
5 x 5 Wiener	0.21	4.68	19.66
9 x 9 Wiener	0.01	3.24	22.86
3 x 3 submedian	-0.02	3.31	22.67
5 x 5 submedian	-0.04	3.88	21.29
double-gated	0.00	0.83	34.69

Finally, for the sake of comparison, the 0.01 probability-of-false-alarm curves are all plotted on the same graph in Figs. 18 and 19. The very same type of lines is used for the two Laplacian filters (solid line), the two Wiener filters (long dash followed by two short dashes), and the two submedian filters (long dash followed by a dot). The double-gated filter is represented by a dashed line.

From this set of figures and tables, we see that:

- a) The filter with the best overall performance for the type of background considered here is the 3 x 3, 8-neighbor Laplacian.
- b) With the exception of the 3 x 3, 8-neighbor Laplacian filter, the double-gated filter outperforms all the other filters at low target intensities (Fig. 18).
- c) The double-gated filter is more robust than all the other filters, without exception, in the presence of noise (Fig. 19).

CLOUD BACKGROUND


($P_f = 0.01$)

FIGURE 18 - We have gathered in this figure the detection-performance curves (at a probability of false alarm of 0.01) of the seven filters studied in this report. The same type of line is used for the two Laplacian filters (solid line), the two Wiener filters (combination of long and short dashes) and the two submedian filters (dash-dot line). A dashed line is used to represent the double-gated filter.

CLOUD BACKGROUND


($P_f = 0.01$)

FIGURE 19 - This figure is the counterpart of Fig. 18 as concerns the behavior of the seven filters in the presence of additive Gaussian noise.




FIGURE 20 - This figure shows the logarithm of the power spectrum of the four linear filters. For display purposes and to bring out the differences between the various spectra, the values of the logarithm between the 0.1 and the 0.9 percentage points were mapped into a gray scale running from 0 (black) to 255 (white), and then subjected to a modulo-64 transformation (i.e. gray levels 0 to 63 were mapped into 0 to 255, as well as the gray levels in the ranges 64 to 127, 128 to 191, and 192 to 255). This last transformation produces a display that looks like an interference pattern.

- d) The Laplacian, Wiener and submedian filters exhibit the same behavior in the presence of noise.
- e) For the cloud background considered in this report, the three nonlinear filters (the double-gated filter and the two submedian filters) do not do better as a whole than the four linear ones.
- f) The curves of Figs. 11 and 13 are nearly identical. This should come as no surprise given that the point-spread functions of the relevant filters are quite similar despite the difference in size. In fact, if [12] is rounded to the first decimal digit, we obtain:

0.0	-0.3	0.0
-0.3	+1.0	-0.3
0.0	-0.3	0.0

Figure 20 further confirms the great similarity between the two filters.

- g) It seems (Figs. 15 and 16) that the performance of a submedian filter decreases as we increase its size.
- h) If the seven filters were to be ranked from best to worst according either to the MDTI (Table I) or the MTNF (Table II) values for a probability of false alarm of 0.01, one would obtain the following classification:

- 1) 3 x 3, 8-neighbor Laplacian
- 2) 3 x 3 submedian filter
- 3) 9 x 9 Wiener filter
- 4) 3 x 3, 4-neighbor Laplacian
- 5) 5 x 5 Wiener filter
- 6) 5 x 5 submedian filter
- 7) 5 x 5 double-gated filter

1) However, if the same sort of ranking was attempted with the AC power gain values of Table III, the result would be:

- 1) 5 x 5 double-gated filter
- 2) 3 x 3, 8-neighbor Laplacian
- 3) 9 x 9 Wiener filter
- 4) 3 x 3 submedian filter
- 5) 5 x 5 submedian filter
- 6) 5 x 5 Wiener filter
- 7) 3 x 3, 4-neighbor Laplacian

This is a clear indication that the ratio of the standard deviation of the filtered image to the one of the unfiltered image is not a general measure of the performance of a spatial filter. It is meaningful only for linear filters.

5.0 CONCLUSION

The primary objective of this report was to formulate a methodology for evaluating both linear and nonlinear spatial filters used by passive IR sensors to discriminate between signals due to target sources and those arising from background clutter. As mentioned in the introduction, the proposed methodology applies only to point targets. A different methodology will have to be devised for spread targets, for it does not seem possible to generalize the one expounded here in order to include such targets.

A number of specific spatial filters cannot be evaluated with the proposed methodology. These are the ones like the Sobel filter whose center element of the point-spread function is zero. This is because we assume that the post-filtering processing is deliberately designed to announce that a detection has occurred if, and only if, the target coincides with the center element of the filter. Otherwise, the target is regarded as being absent.

The proposed methodology allows one to quantitatively and qualitatively evaluate the performance of a particular spatial filter for a given background. The results obtained will be useful only insofar as this particular background is typical of those likely to be encountered by the surveillance system or any other system built around such a filter. In other words, the performance of a spatial filter is bound to be different for different backgrounds.

The results presented in this report, along with those of a few more experiments we have carried out with other types of background, convey the impression that the 3×3 , 8-neighbor Laplacian filter is a very effective filter when the imagery is not too noisy, that the double-gated filter is more robust than the other ones in the presence of noise, and that the performance of a submedian filter decreases as one increases its size. However, more experimentation will be needed to confirm these findings.

6.0 REFERENCES

1. Lithopoulos, E., "Digital Processing for Space Based IR Surveillance", Final Report, Contract No. 05D83-00183, Spar Aerospace Limited, January 1985, UNCLASSIFIED
2. "Infrared Processing Study", Final Report, Contract No. 1SD81-00097, Computing Devices Company, UNCLASSIFIED
3. Kay, I. W., "Evaluation of Infrared Target Discrimination Algorithms", IDA Paper P-1714, April 1983, UNCLASSIFIED
4. Van Trees, H. L., "Detection, Estimation, and Linear Modulation Theory", John Wiley and Sons, Inc., New York, 1968.

<p>DREV R-4421/86 (UNCLASSIFIED)</p> <p>Research and Development Branch, DND, Canada. DREV, P.O. Box 8800, Courcellette, Que. GOA 1R0</p> <p>"On the Characterization of Spatial Filters for Point-Target Detection in IR Imagery" by L. Sévigny</p> <p>This report is a comparative study of the spatial filters proposed for detecting stationary or slow-moving point targets. The filters are characterized by their probability of detection against a) the probability of false alarm, b) the strength or contrast of the target for a given false-alarm rate or c) the noise for a given false-alarm rate and a given strength or contrast.</p>	<p>DREV R-4421/86 (UNCLASSIFIED)</p> <p>Research and Development Branch, DND, Canada. DREV, P.O. Box 8800, Courcellette, Que. GOA 1R0</p> <p>"On the Characterization of Spatial Filters for Point-Target Detection in IR Imagery" by L. Sévigny</p> <p>This report is a comparative study of the spatial filters proposed for detecting stationary or slow-moving point targets. The filters are characterized by their probability of detection against a) the probability of false alarm, b) the strength or contrast of the target for a given false-alarm rate or c) the noise for a given false-alarm rate and a given strength or contrast.</p>
<p>DREV R-4421/86 (UNCLASSIFIED)</p> <p>Research and Development Branch, DND, Canada. DREV, P.O. Box 8800, Courcellette, Que. GOA 1R0</p> <p>"On the Characterization of Spatial Filters for Point-Target Detection in IR Imagery" by L. Sévigny</p> <p>This report is a comparative study of the spatial filters proposed for detecting stationary or slow-moving point targets. The filters are characterized by their probability of detection against a) the probability of false alarm, b) the strength or contrast of the target for a given false-alarm rate or c) the noise for a given false-alarm rate and a given strength or contrast.</p>	<p>DREV R-4421/86 (UNCLASSIFIED)</p> <p>Research and Development Branch, DND, Canada. DREV, P.O. Box 8800, Courcellette, Que. GOA 1R0</p> <p>"On the Characterization of Spatial Filters for Point-Target Detection in IR Imagery" by L. Sévigny</p> <p>This report is a comparative study of the spatial filters proposed for detecting stationary or slow-moving point targets. The filters are characterized by their probability of detection against a) the probability of false alarm, b) the strength or contrast of the target for a given false-alarm rate or c) the noise for a given false-alarm rate and a given strength or contrast.</p>

<p>CRDV R-4421/86 (SANS CLASSIFICATION)</p> <p>Bureau - Recherche et Développement, MDN, Canada. CRDV, C.P. 8800, Courcellette, Qué. GOA 1R0</p> <p>"Caractérisation des filtres spatiaux utilisés pour détecter des cibles ponctuelles en imagerie infrarouge" par L. Sévigny</p> <p>Ce rapport consiste en une étude comparative des différents filtres spatiaux qui sont proposés pour la détection de cibles ponctuelles stationnaires ou quasi stationnaires. Les filtres ont été caractérisés par leur probabilité de détection en fonction: a) de la probabilité d'une fausse alarme, b) de l'intensité ou du contraste de la cible pour une probabilité de fausse alarme donnée ou c) du bruit pour une probabilité de fausse alarme et une intensité ou un contraste donnés.</p>	<p>CRDV R-4421/86 (SANS CLASSIFICATION)</p> <p>Bureau - Recherche et Développement, MDN, Canada. CRDV, C.P. 8800, Courcellette, Qué. GOA 1R0</p> <p>"Caractérisation des filtres spatiaux utilisés pour détecter des cibles ponctuelles en imagerie infrarouge" par L. Sévigny</p> <p>Ce rapport consiste en une étude comparative des différents filtres spatiaux qui sont proposés pour la détection de cibles ponctuelles stationnaires ou quasi stationnaires. Les filtres ont été caractérisés par leur probabilité de détection en fonction: a) de la probabilité d'une fausse alarme, b) de l'intensité ou du contraste de la cible pour une probabilité de fausse alarme donnée ou c) du bruit pour une probabilité de fausse alarme et une intensité ou un contraste donnés.</p>
<p>CRDV R-4421/86 (SANS CLASSIFICATION)</p> <p>Bureau - Recherche et Développement, MDN, Canada. CRDV, C.P. 8800, Courcellette, Qué. GOA 1R0</p> <p>"Caractérisation des filtres spatiaux utilisés pour détecter des cibles ponctuelles en imagerie infrarouge" par L. Sévigny</p> <p>Ce rapport consiste en une étude comparative des différents filtres spatiaux qui sont proposés pour la détection de cibles ponctuelles stationnaires ou quasi stationnaires. Les filtres ont été caractérisés par leur probabilité de détection en fonction: a) de la probabilité d'une fausse alarme, b) de l'intensité ou du contraste de la cible pour une probabilité de fausse alarme donnée ou c) du bruit pour une probabilité de fausse alarme et une intensité ou un contraste donnés.</p>	<p>CRDV R-4421/86 (SANS CLASSIFICATION)</p> <p>Bureau - Recherche et Développement, MDN, Canada. CRDV, C.P. 8800, Courcellette, Qué. GOA 1R0</p> <p>"Caractérisation des filtres spatiaux utilisés pour détecter des cibles ponctuelles en imagerie infrarouge" par L. Sévigny</p> <p>Ce rapport consiste en une étude comparative des différents filtres spatiaux qui sont proposés pour la détection de cibles ponctuelles stationnaires ou quasi stationnaires. Les filtres ont été caractérisés par leur probabilité de détection en fonction: a) de la probabilité d'une fausse alarme, b) de l'intensité ou du contraste de la cible pour une probabilité de fausse alarme donnée ou c) du bruit pour une probabilité de fausse alarme et une intensité ou un contraste donnés.</p>

END

2-87.

DTIC